7. Appendix

Next, we present the Supplementary Materials for the paper "Re-ReND: Real-time Rendering of NeRFs across Devices". Specifically, in addition to the results reported in the paper, we report results of Re-ReND w.r.t. Image Quality (Section 8.1) and (Section 9.1), Rendering Speed (Section 8.2), Mesh Size (Section 8.3 and Section 9.2), Disk Space (Section 8.4), validation of view-dependent effects (Section 10), sensitivity to geometry variations (Section 11) and Photo-metric quality w.r.t. embedding dimensionality D (Section 12). Furthermore, we encourage the reviewers to watch the **associated video**, `Re-ReND.mp4`, demonstrating Re-ReND's capabilities of real-time rendering across devices. This video demonstrates how Re-ReND can render, in real time, a scene composed of tens (Figure 9) or even thousands (Figure 10) of objects. Figure 9 illustrates such a scene, composed of moving chairs, hotdogs, the drumset, and a microphone.

Figure 9. Re-ReND **enables real-time rendering of scenes that can be composed by tens of objects.** Please refer to the video `Re-ReND.mp4`, demonstrating real-time rendering of this scene.

8. Quantitative Results

8.1. Image Quality

At rendering, the image quality achieved by Re-ReND depends on the amount of texels, *i.e.* pixels in the texture map, assigned to each triangle in the mesh. Here, we report the effect that this variable has on image quality. For assessing image quality, we measure the standard quality metrics (PSNR, SSIM [33] and LPIPS [38]) on both datasets (*Realistic Synthetic 360°* and *360° Unbounded Tanks and Temples*). Here we report disaggregate per-scene measures of PSNR (Table 4 and Table 5), SSIM (Table 6 and Table 7), and LPIPS (Table 8 and Table 9) for Re-ReND.

The traditional configuration of texels is equally-sized triangles in a texture map. Thus, the number of texels (*i.e.* the column " Tex." in Tables 4-9) increases quadratically w.r.t. the triangle's side. Formally, $\#\,\text{Tex.} = \lceil p^2/2 \rceil$, where p is the number of pixels in the triangle's side.

Naturally, across all metrics and datasets, image quality improves as the number of texels increases. We find that Re-ReND provides competitive performance when using 18 texels per triangle. However, while performance improves by increasing the number of texels, the quadratic growth of texels makes performance gains rapidly reach diminishing returns.

8.2. Rendering Speed

We report the disaggregated rendering speeds achieved by Re-ReND, measured in frames per second (FPS), in Table 10. For the Realistic Synthetic 360° dataset, Re-ReND attains an average speed of over 54 FPS even on a Samsung S21.

Tex.	chair	drums	ficus	hotd.	lego	mat.	mic	ship	aver.
5	28.69	23.28	27.25	31.71	28.89	26.59	28.69	24.62	27.46
8	29.66	23.75	27.57	32.55	30.13	27.16	29.78	25.04	28.20
13	30.40	24.02	27.73	33.09	30.93	27.49	30.46	25.29	28.68
18	30.99	24.19	27.83	33.48	31.48	27.71	30.91	25.45	29.00
24	31.46	24.31	27.90	33.75	31.87	27.85	31.24	25.56	29.24
32	31.85	24.39	27.94	33.95	32.17	27.95	31.48	25.63	29.42
40	32.17	24.46	27.97	34.09	32.39	28.03	31.66	25.68	29.56
50	32.43	24.50	27.99	34.21	32.57	28.08	31.81	25.72	29.66
60	32.64	24.54	28.01	34.30	32.71	28.13	31.92	25.75	29.75
72	32.83	24.56	28.02	34.37	32.82	28.16	31.91	25.78	29.81

Table 4. **PSNR on the Realistic Synthetic 360° dataset.**

Tex.	truck	train	m60	playg.	average
5	14.92	19.31	16.70	18.53	17.37
8	15.15	19.46	16.87	18.80	17.57
13	15.27	19.56	16.98	18.96	17.69
18	15.34	19.62	17.05	19.06	17.77
24	15.40	19.66	17.10	19.14	17.82
32	15.44	19.69	17.14	19.19	17.87
40	15.47	19.71	17.17	19.23	17.90
50	15.49	19.72	17.20	19.26	17.92
60	15.51	19.74	17.22	19.29	17.94
72	15.52	19.75	17.23	19.31	17.95

Table 5. **PSNR of 360° Unbounded Tanks and Temples dataset.**

Tex.	chair	drums	ficus	hotd.	lego	mat.	mic	ship	aver.
5	0.93	0.9	0.943	0.948	0.933	0.917	0.965	0.807	0.918
8	0.942	0.908	0.947	0.955	0.946	0.926	0.971	0.814	0.926
13	0.95	0.912	0.949	0.959	0.953	0.93	0.975	0.818	0.931
18	0.955	0.915	0.95	0.962	0.958	0.933	0.977	0.821	0.934
24	0.959	0.917	0.951	0.964	0.961	0.935	0.978	0.822	0.936
32	0.963	0.918	0.951	0.965	0.963	0.936	0.979	0.824	0.937
40	0.965	0.919	0.952	0.966	0.964	0.937	0.98	0.824	0.938
50	0.967	0.92	0.952	0.967	0.965	0.938	0.98	0.825	0.939
60	0.968	0.92	0.952	0.967	0.966	0.938	0.98	0.825	0.940
72	0.969	0.921	0.952	0.968	0.966	0.939	0.98	0.826	0.940

Table 6. **SSIM of Realistic Synthetic 360° dataset.**

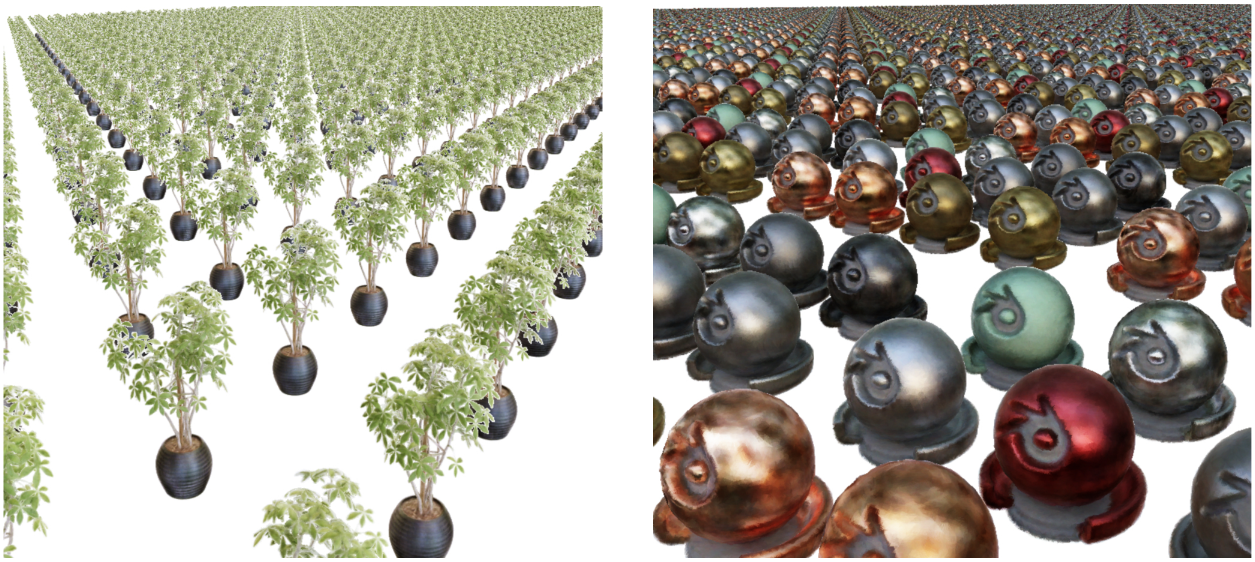

Figure 10. **With Re-ReND, we can simultaneously render thousands of objects in real time.** Here we show 2,500 ficus scenes (left) and materials scenes (right). Please refer to the accompanying video `Re-ReND.mp4`, demonstrating real-time rendering of these scenes.

Tex.	truck	train	m60	playg.	average
5	0.486	0.513	0.456	0.535	0.498
8	0.498	0.524	0.468	0.551	0.510
13	0.506	0.533	0.477	0.563	0.520
18	0.514	0.538	0.484	0.571	0.527
24	0.521	0.543	0.489	0.578	0.533
32	0.525	0.546	0.493	0.583	0.537
40	0.529	0.548	0.496	0.586	0.540
50	0.532	0.55	0.498	0.589	0.542
60	0.535	0.551	0.5	0.591	0.544
72	0.537	0.552	0.501	0.593	0.546

Table 7. **SSIM of 360° Unbounded Tanks and Temples dataset.**

Tex.	chair	drums	ficus	hotd.	lego	mat.	mic	ship	aver.
5	0.072	0.119	0.069	0.093	0.102	0.105	0.062	0.219	0.105
8	0.062	0.109	0.061	0.079	0.083	0.094	0.052	0.209	0.094
13	0.055	0.101	0.056	0.07	0.069	0.086	0.046	0.202	0.086
18	0.049	0.094	0.053	0.064	0.06	0.08	0.041	0.197	0.080
24	0.045	0.09	0.051	0.059	0.053	0.076	0.037	0.193	0.076
32	0.042	0.086	0.049	0.056	0.049	0.073	0.035	0.191	0.073
40	0.039	0.084	0.048	0.053	0.045	0.07	0.032	0.189	0.070
50	0.037	0.082	0.047	0.052	0.042	0.068	0.031	0.187	0.068
60	0.035	0.08	0.046	0.051	0.04	0.066	0.029	0.186	0.067
72	0.034	0.079	0.046	0.05	0.038	0.065	0.028	0.185	0.066

Table 8. **LPIPS of Realistic Synthetic 360° dataset.**

Tex.	truck	train	m60	playg.	average
5	0.526	0.522	0.57	0.517	0.534
8	0.519	0.514	0.565	0.507	0.526
13	0.515	0.512	0.564	0.501	0.523
18	0.513	0.51	0.563	0.496	0.521
24	0.511	0.509	0.562	0.492	0.519
32	0.509	0.507	0.56	0.49	0.517
40	0.508	0.506	0.559	0.487	0.515
50	0.506	0.505	0.557	0.485	0.513
60	0.505	0.504	0.555	0.483	0.512
72	0.504	0.504	0.553	0.481	0.51

Table 9. **LPIPS of 360° Unbounded Tanks and Temples dataset.**

8.3. Mesh Size

Table 4 in the main paper reports the average size (number of vertices and triangle faces) of the meshes used by Re-ReND. We report the per-scene mesh sizes in Table 11 for both datasets we experimented with. Overall, for the Realistic Synthetic 360° dataset (left columns), Re-ReND uses, on average, fewer than 205k faces and 99k vertices. For the 360° Unbounded Tanks and Temples dataset, these numbers correspond to 250k faces and 120k vertices. As such, these meshes are decidedly not particularly precise, and thus serve mostly as a collision mesh for Re-ReND to estimate where the scene's geometry is.

8.4. Disk Space

The number of texels assigned to each triangle in the mesh affects the disk space used for representing a scene. We vary the number of texels, and report the disk space used for each scene in Table 12 (for Realistic Synthetic 360°) and Table 13 (for 360° Unbounded Tanks and Temples).

For the Realistic Synthetic 360° dataset, Re-ReND's default of 18 texels implies using an average disk space of 198.8 MB. Furthermore, all the objects (except ficus, lego and ship), use fewer than 200 MB. On the other hand, for the 360° Unbounded Tanks and Temples dataset, the default of 18 texels makes all scenes use a disk space between 270 and 310 MB.

9. Qualitative Results

9.1. Image Quality

In Figure 12, and Figure 13, we present the qualitative results obtained on a synthetic dataset. The images in the first

	chair	drums	ficus	hotdog	lego	materials	mic	ship	aver.	truck	train	playground	m60	aver.
Samsung S21	60.1	60.0	60.1	60.1	36.8	60.1	60.0	40.7	54.7	23.2	37.0	27.9	45.7	33.5
Motorola G9	11.7	13.4	6.5	12.3	4.1	14.2	18.0	3.2	10.4	6.7	9.0	5.8	9.4	7.7
Galaxy S6	31.6	33.8	21.1	27.2	12.5	32.3	42.5	12.2	26.6	18.1	20.7	19.6	23.3	20.4
Dell	84.6	84.5	74.8	72.4	50.3	82.3	110.4	42.8	75.3	49.8	57.1	54.4	54.8	54.0
Gaming	769.2	762.7	688.3	684.4	447.2	759.5	1065.1	401.8	697.3	469.4	560.8	483.2	553.0	516.6
PC	1113.1	1130.7	997.0	1067.6	807.6	998.9	1117.0	873.4	1013.2	884.3	967.0	952.1	898.2	925.4

Table 10. **Frames per second (FPS) achieved by** Re-ReND**.** We report the disaggregated FPS for all devices we tested on all the scenes. Columns 2-10: Realistic Synthetic 360° dataset. Columns 11-15: 360° Unbounded Tanks and Temples dataset.

	chair	drums	ficus	hotdog	lego	materials	mic	ship	average	truck	train	playground	tank	average
Faces	158k	164k	239k	146k	360k	159k	131k	284k	205,693	265k	232k	246k	235k	244,847
Vertices	76k	80k	119k	72k	171k	76k	62k	136k	99,539	125k	115k	116k	115k	117,751

Table 11. **Mesh sizes used by** Re-ReND**.** We report the number of triangle faces and vertices used to model each scene. Columns 2-10: Realistic Synthetic 360° dataset. Columns 11-15: 360° Unbounded Tanks and Temples dataset.

# Tex.	chair	drums	ficus	hotd.	lego	mater.	mic	ship	average
5	69.1	75.7	107.3	66.4	141.7	73.4	60.5	111.3	88.1
8	94.8	103.4	147.7	88.7	195.6	98.7	80.9	149.4	119.8
13	126.0	135.7	193.9	114.4	259.1	128.0	105.1	193.3	156.9
18	161.8	172.2	245.9	143.0	330.1	160.8	132.3	241.8	198.4
24	201.7	212.0	303.0	174.0	408.3	196.7	162.6	294.7	244.1
32	245.8	255.7	364.5	207.8	492.7	235.5	195.8	351.1	293.6
40	293.6	302.1	430.0	243.5	583.0	276.9	231.2	411.2	346.4
50	344.7	351.5	499.2	281.2	678.8	320.6	269.2	474.2	402.4
60	399.2	403.6	571.7	320.6	779.6	366.5	309.4	540.3	461.3
72	456.9	458.0	647.3	361.8	885.0	414.6	349.7	609.1	522.8

Table 12. **Disk Space (MB) of Realistic Synthetic 360° dataset.**

# Tex.	truck	train	m60	playg.	average
5	143.8	132.3	147.6	137.7	140.37
8	185.3	172.3	192.9	179.1	182.42
13	234.0	220.1	246.3	228.3	232.17
18	288.7	273.6	307.0	283.9	288.31
24	348.3	333.1	374.4	345.7	350.37
32	413.3	398.1	448.2	412.5	418.04
40	482.1	468.0	527.6	484.3	490.49
50	555.1	542.3	612.7	561.5	567.89
60	631.4	620.7	702.7	642.3	649.27
72	712.7	702.6	798.1	727.4	735.18

Table 13. **Disk Space (MB) of 360° Unbounded Tanks and Temples dataset.**

column represent the ground truth (GT) data. In the second column, we show the results obtained using Re-ReND with quad size 72 and 32 components. Finally, in the third column, we display the results obtained using Re-ReND with quad size 18 and 32 components.

Upon visual inspection, we observe that Re-ReND with quad size 72 and 32 components produces more accurate and visually appealing results compared to Re-ReND with quad size 18 and 32 components. The former shows greater detail and smoother transitions between the different regions of the scene. However, Re-ReND with quad size 18 and 32 components still manages to produce decent results.

In addition to the synthetic dataset, we also present qualitative results on a real 360 dataset [See Figure 14]. The images in the first column represent the ground truth (GT) data. In the second column, Re-ReND before discretization, while in the third column, Re-ReND with quad size 18 and 32 components.

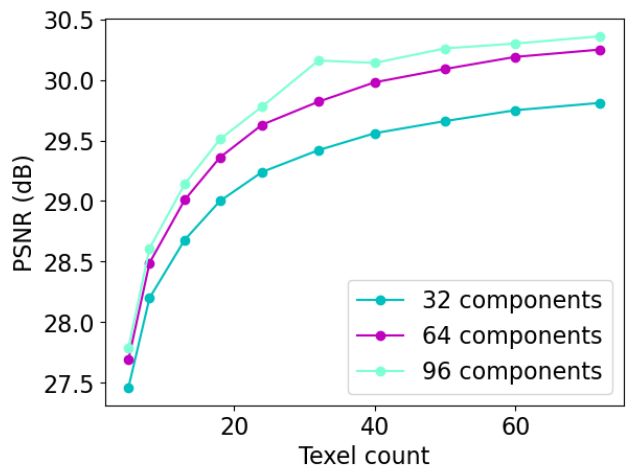

Figure 11. **Photo-metric quality depending on the dimensionality** We report Re-ReND's results for various dimensionalities of embedding.

It is worth noting that these qualitative results are obtained on a real 360 dataset, which presents more challenges compared to the synthetic dataset. The real-world scenario involves more complex lighting conditions, occlusions, and variations in scene geometry.

9.2. Meshes

In Figure 15, we report the meshes we use in Re-ReND for both datasets. All our meshes are simple and smooth. Note that, for the 360° Unbounded Tanks and Temples dataset (last row in Figure 15), the scene is encapsulated within a semi-sphere and a plane mimicking the floor.

10. Validation of view-dependent effects

To validate the effect of view direction, we conducted an experiment comparing the performance of Re-ReND to a simple RGB textured mesh representation as a baseline. Due to the lack of ground truth RGB textures, we created a texture by assigning colors based on the intersected face's normal of a pretrained Re-ReND. The results showed that on both the Synthetic and Unbounded T&T datasets, the RGB textured mesh PSNR was lower compared to Re-ReND. Specifically, the PSNR values were 22.82 dB and 14.79 dB for the RGB textured mesh representation, compared to Re-ReND's 29.00 dB and 17.77 dB, respectively.

This significant performance drop highlights the critical importance of modeling view-dependent effects for achieving high-quality image reconstruction. The Re-ReND approach, which takes into account view-dependent effects, was able to produce more accurate and visually appealing results than the simple RGB textured mesh representation. This finding suggests that the Re-ReND method is effective in modeling view-dependent effects and can lead to improved image reconstruction results.

11. Sensitivity to geometry variations

We conduct an experiment to evaluate the sensitivity of a Re-ReND to geometry quality. To do so, we use the ground truth meshes of the synthetic dataset to train Re-ReND, and we compared two sets of results: one using perfect geometry, and another using "cheap" meshes by marching cubes.

The results were then evaluated using three metrics before discretization: PSNR, SSIM, and LPIPS. The results for the perfect geometry case were 31.10, 0.954, and 0.0535 for PSNR, SSIM, and LPIPS, respectively. For the cheap mesh case, the results were 30.73, 0.946, and 0.0562 for the same metrics.

Analyzing the results, it appears that the Re-ReND is sensitive to geometry quality, as the results for the perfect geometry case were consistently better across all three metrics. This suggests that the method performs better when it has access to high-quality geometry information. However, even when using cheaper meshes, the Re-ReND method can still perform reasonably well. For example, the PSNR values were only slightly different between the perfect and cheap mesh cases, and the difference in LPIPS and SSIM values was within a reasonable range. This suggests that the Re-ReND method can perform very well without losing too much quality, even when the geometry information is not perfect.

12. Photo-metric quality depending on the dimensionality D

In Figure 11, we report Re-ReND's results for various dimensionalities of embedding. The dimensionality of the embedding is an important factor that can affect the performance of the model in various ways. On one hand, a higher dimensional embedding can potentially capture more complex textures and materials, leading to better PSNR. On the other hand, a higher dimensional embedding may also require more memory usage, making it slower and not apt for certain devices.

In practice, the choice of embedding dimensionality is often a trade-off between quality and efficiency, and depends on the specific requirements and constraints of the application. For example, in low-constraint devices, a lower dimensional embedding may be sufficient to achieve good performance, while for desktop, a higher dimensional embedding may be necessary to obtain better results in 8K resolution.

Figure 12. **Qualitative results.** We report the qualitative results for synthetic dataset. From left to right: first column is GT, second column is Re-ReND using quad size 72 and 32 components, and third columns is Re-ReND using quad size 18 and 32 components.

Figure 13. **Qualitative results.** We report the qualitative results for synthetic dataset. From left to right: first column is GT, second column is Re-ReND using quad size 72 and 32 components, and third columns is Re-ReND using quad size 18 and 32 components.

Figure 14. **Qualitative results.** We report the qualitative results for real 360 dataset. From left to right: first column is GT, second column is Re-ReND using before discretization, and third columns is Re-ReND using quad size 18 and 32 components.

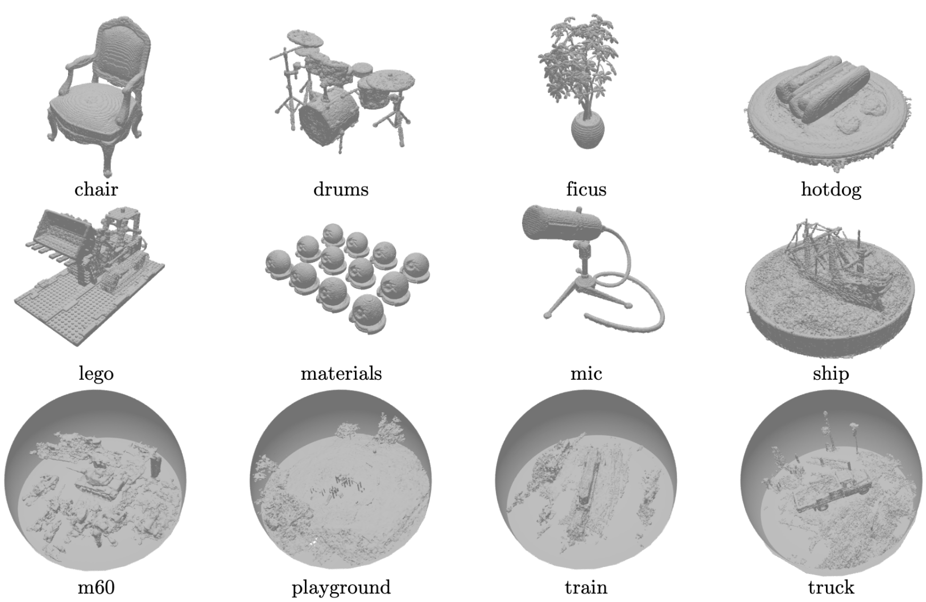

Figure 15. **Meshes used by** Re-ReND**.**